arXiv:1702.08001v1 [cs.LG] 26 Feb 2017

Bayesian Nonparametric Feature and Policy Learning for Decision-Making

Jürgen Hahn* Abdelhak M. Zoubir†
Signal Processing Group
Institute of Telecommunications, Technische Universität Darmstadt
Merckstraße 25, 64283 Darmstadt

Abstract

Learning from demonstrations has gained increasing interest in the recent past, enabling an agent to learn how to make decisions by observing an experienced teacher. While many approaches have been proposed to solve this problem, there is only little work that focuses on reasoning about the observed behavior. We assume that, in many practical problems, an agent makes its decision based on latent features, indicating a certain action. Therefore, we propose a generative model for the states and actions. Inference reveals the number of features, the features, and the policies, allowing us to learn and to analyze the underlying structure of the observed behavior. Further, our approach enables prediction of actions for new states. Simulations are used to assess the performance of the algorithm based upon this model. Moreover, the problem of learning a driver's behavior is investigated, demonstrating the performance of the proposed model in a real-world scenario.

Key words: Bayesian nonparametrics, decision-making, learning from demonstrations, feature learning, imitation learning

1 Introduction

Decision-making plays a crucial role in many applications, such as robot learning, driver assistance systems, and recommender systems [40]. A fundamental question in decision-making is how an agent can learn to make optimal decisions. Learning from an experienced teacher provides a natural means to solve this problem, without the need of explicitly defining rules for the desired behavior. Further, observing a teacher may provide a deeper understanding of the decision-making process. Therefore, Learning From Demonstrations (LFD) [1] has gained a lot of interest in the recent past.

*jhahn@spg.tu-darmstadt.de
†zoubir@spg.tu-darmstadt.de

According to [1], approaches for LFD can be grouped into (i) reward-based models and (ii) imitation learning. In reward-based models, it is assumed that the agent makes its decision based on a reward which is, in the context of LFD, learned from observations (as in Inverse Reinforcement Learning (IRL) [30]). In imitation learning, it is assumed that an experienced teacher can be observed. Thus, the policy, telling the agent how to act in a given situation, can be learned by understanding the direct relation between the teacher's states and actions. Especially in reward-based approaches, problems defined on infinite spaces, where the state of an agent can take continuous values, are mostly intractable to solve and efficient approximations are needed. Only in the case of finite state and action spaces and known rewards, learning optimal policies has been solved [3]. A typical approach to facilitate this problem is the representation of the state space in a feature domain, c.f. [33, 5].

However, decision-making has been viewed only from a limited feature-based perspective, where features are usually designed by experts and mainly serve the purpose of reducing the size of the state space. We argue that, in many practical systems, the agent makes its decision based on a compact representation of the observed data, which can be considered as a projection onto a feature space of the decision-making problem. As an example, consider a person driving a car. The driver's observations consists of, i.a., the location, speed, and acceleration of his vehicle and the surrounding vehicles, the type of the road and the weather conditions. However, the driver makes his decision based on a subset of the available information, e.g., on the time-to-reach between his and the other road users' cars. This idea has also been investigated from a psychological point of view by the concept of discovering latent causes in human behavior, which is related to learning state space representations [12].

Assuming that a certain structure underlies the observations, we aim at inferring the latent features and build a feature-based representation of the states, yielding the following two advantages: first, the states can be represented compactly, rendering the decision-making problem much more efficient as compared to working in the original domain. Second, the features can be regarded as causes for the observed decisions, allowing for a deeper understanding of the observed behavior. In particular, we consider a LFD problem and propose a Bayesian nonparametric framework for feature learning in LFD. We assume that the policies depend on the features of the observed demonstrations. With this model, we are able to (i) significantly reduce the state space by (ii) learning the features as well as the number of features and (iii) provide a better understanding of what caused the teacher to take the observed actions.

In the next parts of this section, we formulate the problem and review the state-of-the-art for feature learning in LFD. In Section 2, we motivate and present the model for feature-based decision-making. The model is detailed in Section 3 and we explain how the number of latent features can be inferred from the data. Section 4 presents the inference scheme for the latent variables. The proposed algorithm is empirically evaluated with simulation experiments in Section 5, demonstrating its performance. Real data experiments are conducted in Section 6. We discuss our findings in Section 7 and end with a short conclusion in Section 8.

1.1 Problem formulation

The goal of this work is to introduce a feature-based LFD framework. While this model can be used for teaching an agent given observations of the desired behavior, the focus of this work lies on the analysis of the observed behavior by means of the features. Since we consider a decision-making task, the investigated problem can be modeled by means of a Partially Observable Markov Decision Process (POMDP) [21, 42], which is defined by

- a set of observations, $\mathcal{Z}$,
- a set of states, $\mathcal{X}$,
- a finite set of N_u actions, $\mathcal{U}$,
- a transition model, which describes the probability of entering a state after taking an action in the current state,
- an observation model which explains how the observations are generated from the states,
- a discount factor, which penalizes long-term rewards,
- and a reward function, R.

As the main goal of this work is to provide a means to understand observed behavior, we consider an imitation learning approach and learn the relation between states and actions from the observations directly, where the reward and, hence, the discount factor, are not considered.

We assume that we (or an agent) have access to N_z noisy observations, $\mathbf{z}_n \in \mathcal{Z}$, of the states, $\mathbf{x}_n \in \mathcal{X}$, with $n = 1, \ldots, N_\mathrm{z}$. In particular, we consider Gaussian noise, i.e., $\mathbf{z}_n$, describes observations of the states, $\mathbf{x}_n$, with additive Gaussian noise. Further, we assume that the actions, $u_n \in \mathcal{U}$, taken in the corresponding states, can be observed. The observations are assumed to be optimal in the sense that they represent the behavior of the agent seeking its goal, i.e., without any exploratory steps.

A simple approach to the considered problem would be the use of a feature extraction technique such as Principal Component Analysis (PCA) [20] or Non-negative Matrix Factorization (NMF) [23] to learn features from the observed states. Following this solution, features cannot be jointly learned and shared between different actions. Moreover, learning discriminative features is not guaranteed. Therefore, we argue that the features and policies need to be learned jointly such that a trade-off between feature sharing, promoting a compact model, and discrimination capability is found based on the observations.

1.2 Relation to existing work

In the following, we provide an overview about the current state of the art of feature learning for LFD. As IRL usually requires to solve an Markov Decision Process (MDP), which is usually done by means of a Reinforcement Learning (RL) algorithm, we start with an introduction in feature learning for RL.

1.2.1 Reinforcement Learning

Large state and action spaces are especially problematic for value-based RL algorithms such as value-iteration [3] or Q-learning [48] since the value function, representing the expected accumulated reward, needs to be approximated. In early approaches, a set of basis functions, often referred to as features, is linearly weighted to represent the this function [5]. However, there exists only little work on learning these basis functions. Riedmiller [38] has proposed the Q-fitted value iteration where the value function is approximated by means of a neural network, where the features are learned in the layers of the network. In [29], this approach has been extended by replacing the neural network by a deep-layered counterpart. A different concept to employ features is proposed by Hutter with the Feature Reinforcement Learning framework [18]. In this framework, the goal is to learn a feature mapping from the agent's history (comprised of actions, states, and rewards) to a MDP state, enabling decision learning for infinite state spaces [8]. An alternative framework for learning the latent structure of the state space is proposed in [9] which is based on Factored Markov Decision Processs (FMDPs) [4]. In the FMDP framework, it is assumed that the observed states can be represented compactly by exploiting the structure within the states, enabling efficient learning. This approach has been extended in [31] to an online approach, where the features are selected from a large set by means of Group LASSO. In these approaches, the inferred features mainly serve the purpose of dimensionality reduction. Therefore, the features do not necessarily possess a meaning that can be easily interpreted.

1.2.2 Inverse Reinforcement Learning

IRL is concerned with the problem of learning the reward function from observed behavior [30]. As in RL, features are often used to linearly parameterize the reward function, e.g., in [30, 17]. Recent attempts have been made to consider a Deep Learning (DL) architecture [49] for IRL, providing means for nonlinear, hierarchical feature learning.

A Bayesian nonparametric approach is proposed in [6], utilizing an Indian Buffet Process (IBP) to model feature activations. As the features of the reward function are assumed to be known, this approach can be understood rather as a feature selection than feature learning for IRL, where the number of features is inferred by the IBP. Different results on Bayesian nonparametrics for IRL, which is indirectly related to feature learning, are given in [27], where a partitioning of the state space is sought for, or [45], where complex behavior is decomposed into several, simpler behaviors that can be easily learned.

1.2.3 Imitation Learning

Instead of estimating the reward as in IRL, imitation learning aims at inferring the underlying policy directly [2, 44]. Usually, handcrafted features are used, e.g., in [37, 39]. Attempts to introduce new features are made in [36] as an extension of the maximum margin planning algorithm which is proposed in [37]. As explained in [1], imitation learning can be considered

as a supervised learning task. Thus, feature selection and learning techniques developed for classification and regression can also be used in imitation learning. An excellent overview is given in [16]. Although these models work well in practice, they might not be able to provide a deeper understanding of the observed behavior as they do not explicitly model the states and actions.

2 Choice of the Model

In the first part of this section, we propose a feature model for LFD. In the second part, we explain the relevance of the transition model to the proposed framework. Since we assume a mixture of policies in this framework, we briefly discuss the intuition behind this assumption in the third part. Alternative models for feature learning for LFD are discussed in the forth part.

2.1 Feature model for learning from demonstrations

We assume a linear latent feature model, similar to NMF [23] and PCA [20]. Thus, the noisy observations, $\mathbf{z}_n \in \mathbb{R}^{1\times D}, n = 1, \ldots, N_z$, are assumed to be composed of the latent features, $\mathbf{F} \in \mathcal{F}^{K\times D}$, and the feature coefficients, $\mathbf{s}_n \in \mathcal{S}^{1\times K}$,

$$\mathbf{z}_n = \mathbf{s}_n\mathbf{F} + \boldsymbol{\epsilon}_n, \tag{1}$$

where $\boldsymbol{\epsilon}_n$ represents Gaussian i.i.d. noise with variance σ_z^2 and the states are given as $\mathbf{x}_n = \mathbf{s}_n\mathbf{F}$. The number of features is K and the dimension of the observations is D. Clearly, the feature space, $\mathcal{F}$, depends on the application. In the following, we assume that the features are positive-valued, i.e., $\mathcal{F} = \mathbb{R}_+$. Following [22], the feature matrix is composed of a binary activation matrix, $\mathbf{A} \in \{0, 1\}^{K\times D}$, and a weighting matrix, $\mathbf{W} \in \mathcal{F}^{K\times D}$, where the relation is given by the Hadamard product,

$$\mathbf{F} = \mathbf{A} \odot \mathbf{W}. \tag{2}$$

The feature model in Eq. (2) can be easily extended to an infinite feature model by placing an IBP [13, 15] prior over $\mathbf{A}$. The IBP assumes an infinite number of features, while the observed data can be explained by a finite number. This gives rise to a nonparametric model, where the number of features is implicitly modeled by means of the IBP. This is detailed in Section 3.

A fundamental difference to most existing work on decision-making using feature representations is that we assume that the agent makes its decision based on features, where each feature attracts the agent to take a specific action. We consider a (latent) linear feature model where the feature coefficients depend on the observed state and determine the actions. For this reason, we refer to the feature coefficients also as *substates*.

2.2 Transition model

In decision-making problems, the transition model explains which actions the agent can take in each state. Thus, the transition model acts as a constraint on the possible actions. As we focus on inferring the latent causes for the observed actions by means of features, the transition model is less relevant as it only provides additional information. Besides, the transition model is rarely known in practice. As we are given observations, we can, theoretically, infer the transition model. For this, we could either employ a parametric or a nonparametric model. Defining a parametric model for the transitions is not trivial, as the dynamics in the latent space may be highly nonlinear. Alternatively, a nonparametric model can be assumed. This, however, requires a large amount of observations for the estimation of the parameters, which is often not available. Assuming that the noise in the observations is low, we argue that we can reliably infer the substates from the corresponding observations, eliminating the need for a transition model.

2.3 Feature-based policy

As each feature imposes its on own policy, the probability of the agent taking an action, u, in a substate, $\mathbf{s}$, is a mixture of the feature policies, $P(u \,|\, \boldsymbol{\phi}_k)$,

$$P(u \,|\, \mathbf{s}, \boldsymbol{\Phi}) \propto \sum_{k=1}^{K} s_k P(u \,|\, \boldsymbol{\phi}_k), \tag{3}$$

where $\boldsymbol{\phi}_k$, $k = 1, \ldots, K$, are the parameters of the feature policies and $\boldsymbol{\Phi} = [\boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_K]$. The mixture of policies can be interpreted either as a stochastic or a deterministic policy. In the first case, the action to be taken should be sampled according to Eq. (3). In the second case, simply the most probable action is taken. Mixture polices have been investigated in multi-objective problems, where an agent aims at reaching several objectives, some of which can even be conflicting [32, 46, 41]. A stochastic policy is needed in this framework to ensure that all goals can be satisfied. Similar problems occur if the problem at hand is described by a POMDP, where the true state of the agent is unknown. The uncertainty about the state of the agent is expressed by beliefs over states. Acting according to a stochastic policy maximizes the expected return [21].

In case of single agents, it has been shown that deterministic policies are optimal solutions for MDPs [34]. As explained in Section 2.1, we argue that we can reliably infer the substates. Thus, we assume a deterministic policy in the following, where the probability in Eq. (3) expresses our confidence about the actions given the states.

Note that we also assume that each feature imposes a deterministic policy such that we expect a strongly peaked distribution for the feature policies, expressing the confidence of the chosen action.

2.4 Alternative Feature-based Models

Of course, there are other possibilities to incorporate feature learning in LFD. We briefly discuss two alternatives with their potential advantages and drawbacks. For completeness, as mentioned in the introduction, in the most trivial setup, one could simply cluster the states according to the observed actions and then learn the features. This, however, has the significant drawback that the features cannot be shared between the clusters.

2.4.1 Unique coefficient model

Instead of assuming, as in our model, that the features determine the behavior of the agent, the feature coefficients can be clustered, where the clusters indicate the optimal actions given the coefficients. Thus, the clusters can be interpreted as latent substates. Supposing that the elements of the feature coefficients are binary-valued, we can convert each feature coefficient vector to a unique identifier, representing the cluster. Thus, instead of employing expensive clustering, a fast deterministic mapping from the binary coefficients to the cluster identifier can be used. However, as the identifiers must be unique for the cluster assignments, we can have at maximum 2^K different clusters in this setup, i.e., the number of clusters (and, hence, possible actions) is strictly limited by the number of features. An advantage of this model is that the relation between substates and policies can be nonlinear. However, a significant drawback is that this model suffers severely from errors in the inference of the substates. Consider the case, where, for instance, due to noise, one element of the substate is incorrectly set. This substate is then assigned to a new cluster, for which the policy must be inferred from potentially little data, yielding highly varying policy estimates. As in our approach, an IBP prior can be placed over the feature activations to infer the number of latent features.

2.4.2 Clustering-based approach

A clustering-based approach assumes that similar states can be grouped and result in the same behavior. This can also be understood as a single feature model, in which we assume that the observations can be described by a single feature, representing the clusters. Thus, the substates reduce to cluster indicators, i.e., they indicate which feature best represents the observed state. The number of latent states is then equal to the number of features. One possibility to infer this number is to utilize a Chinese Restaurant Process (CRP), giving rise to a Bayesian nonparametric model [27]. A similar model, where the state space is clustered according to the played actions, is proposed in [43].

2.4.3 Relation to the proposed model

Our model has the advantage, as opposed to the clustering-based approach and the unique coefficient model, that the features provide a means to understand the observed behavior. In contrast to the unique coefficient model, we neither require binary coefficients nor a clustering

step. Compared to the clustering-based approach, our model is able to significantly reduce the number of latent features and policies, as the features can be shared by different states. However, our model suffers from the assumption of a linear relation between features, substates, and policies. This problem is alleviated in the other approaches, as the features are decoupled from the policies. Note that our model becomes similar to the clustering-based model at the price of higher computational costs, if each substate is represented by only one feature.

3 Bayesian Nonparametric Model for Feature Learning

In this section, we provide a general framework for Bayesian nonparametric feature learning for decision-making based on the model proposed in Section 2.1. In order to learn the structure, we assume that we are given a set of observations consisting of state-action pairs, $\mathcal{D} = \{(\mathbf{z}_1, u_1), \ldots, (\mathbf{z}_{N_\mathrm{z}}, u_{N_\mathrm{z}})\}$.

3.1 Observation likelihood and noise variance

The observations, $\mathbf{z}_n, n = 1, \ldots, N_\mathrm{z}$, are assumed to be conditionally independent. As we assume Gaussian noise in Eq. (1), the state likelihood can be expressed as

$$p(\mathbf{Z} \,|\, \mathbf{W}, \mathbf{A}, \mathbf{S}, \sigma_\mathrm{z}^2) = \prod_{n=1}^{N_\mathrm{z}} \mathcal{N}_{\mathbf{z}_n}\left(\mathbf{s}_n \left(\mathbf{A} \odot \mathbf{W}\right), \sigma_\mathrm{z}^2 \mathbf{I}\right) . \tag{4}$$

with $\mathbf{Z} = \begin{bmatrix} \mathbf{z}_1{}^\mathrm{T} & \ldots & \mathbf{z}_{N_\mathrm{z}}{}^\mathrm{T} \end{bmatrix}^\mathrm{T}$ and $\mathbf{S} = \begin{bmatrix} \mathbf{s}_1{}^\mathrm{T} & \ldots & \mathbf{s}_{N_\mathrm{z}}{}^\mathrm{T} \end{bmatrix}^\mathrm{T}$. The variance of the noise, σ_z^2, is assumed to be Inverse-Gamma distributed with hyperparameters $\alpha_\sigma, \beta_\sigma$. Further, we place hyperpriors on α_σ and β_σ, following Gamma distributions with hyperparameters $h_{\alpha_\sigma}^{(1)}, h_{\alpha_\sigma}^{(2)}, h_{\beta_\sigma}^{(1)}$, and $h_{\beta_\sigma}^{(2)}$.

3.2 Prior for the feature weights

As explained above, the prior probability of the feature weights, $\mathbf{W}$, depends on the problem at hand. We consider i.i.d. positive-valued feature weights, $w_{k,d}$ with $k = 1, \ldots, K$ and $d = 1, \ldots, D$, and assume an Exponential prior,

$$p(\mathbf{W} \,|\, \gamma_\mathrm{w}) = \prod_{k=1}^{K} \prod_{d=1}^{D} \mathrm{Exp}_{w_{k,d}}(\gamma_\mathrm{w}) .$$

The scaling factor, γ_w, is assumed to be Inverse-Gamma distributed with hyperparameters α_γ and β_γ.

If the features are assumed to be real-valued, the prior can be modeled with a Gaussian distribution with straightforward modifications.

3.3 Prior for the feature activations

The feature activations are modeled by means of an IBP [13, 15], assuming an infinite number of features. In the following, we consider the two-parameter generalization [14] which allows to sample sparse as well as dense matrices.

The IBP is derived as follows. For a finite number of features, $K^\star$, the sums over the rows of the feature activation matrix, $\mathbf{A}^\star \in \{0,1\}^{K^\star \times D}$, are assumed to follow i.i.d. Binomial distributions. Placing a Beta prior with hyperparameters $\frac{\alpha_a \beta_a}{K^\star}$ and β_a over the parameter of the Binomial distribution and marginalizing over this parameter yields a Beta-Binomial distribution [13, 14]. Since we are interested in sampling from an infinite number of features, we consider the limit for $K^\star \to \infty$, resulting in the distribution of the activation matrix $\mathbf{A}$ [13, 14],

$$\begin{aligned} P(\mathbf{A} \,|\, \alpha_a, \beta_a) &\coloneqq \lim_{K^\star \to \infty} P(\mathbf{A}^\star \,|\, \alpha_a, \beta_a) \\ &= \frac{(\alpha_a \beta_a)^K}{\prod_{\mathbf{h} \in \{0,1\}^D \setminus \mathbf{0}} K_h!} \exp\{-\bar{K}\} \prod_{k=1}^{K} \mathrm{B}(m_k, D - m_k + \beta_a) \ , \end{aligned} \tag{5}$$

where the number of occurrences of the binary vector $\mathbf{h} \in \{0,1\}^D$ is denoted by K_h and B is the Beta function. Note that $\mathbf{A}$ is a matrix with infinitely many rows. However, due to the sparsity assumption, only a finite number of the rows of the realizations contain active elements, denoted by K. Thus, we need to store only rows with active elements in memory, which can be understood as realizations of K features, where the average number of active rows is given by $\bar{K} = \alpha_a \sum_{d=1}^{D} \frac{\beta_a}{\beta_a + d - 1}$ [13, 10]. The hyperparameters α_a and β_a reflect our prior belief about the number of features and the sparsity of the matrix [14]. The probability of activating an element increases with both hyperparameters. As the number of expected features grows linearly with α_a, α_a can be used to control the number of generated features. Considering the limits for $\beta_a \to 0$ and $\beta_a \to \infty$ shows that the expected number of features is limited to $\alpha_a D$. As the probability of an element of $\mathbf{A}$ being active increases, β_a can be understood as a means to control the sparsity of the realizations. We infer α_a and β_a from the observations, assuming that they follow Gamma distributions, i.e., $\alpha_a \sim \mathrm{Ga}_{\alpha_a}\left(h^{(1)}_{\alpha_A}, h^{(2)}_{\alpha_A}\right)$ and $\beta_a \sim \mathrm{Ga}_{\beta_a}\left(h^{(1)}_{\beta_A}, h^{(2)}_{\beta_A}\right)$ [22].

3.4 Prior for the substates

As formulated in Eq. (1), we assume that the observations are composed of a mixture of features weighted by the substates. Similar to a FMDP[1], we restrict the domain of the substate elements, $s_{n,k}$, with $n = 1, \ldots, N_\mathrm{z}$ and $k = 1, \ldots, K$, to take values from a finite set, $\mathcal{S} = \{\breve{s}_1, \ldots \breve{s}_L\}$, where L denotes the number of elements, to simplify inference. For convenience, we assume equidistant elements in $\mathcal{S}$ and set $\breve{s}_1 = 0$ and $\breve{s}_L = 1$. Note that the limited range does not restrict the model, as the features can be scaled to fit the observations.

Further, we assume that the substates are sparse, meaning that each observation consists of

[1] A fundamental difference between our model and a FMDP is that the substates in our model are, theoretically, not restricted to be finite as we do not need to enumerate over them.

only a few features such that only few polices determine the observed action. In particular, we consider a sparsity-promoting mixture prior on the substates, similar to a *spike* and *slab* model [28, 19]. The components are given by a Categorical distribution, where $P(s = 0 \,|\, \theta_{s=0})$ is the sparsity component and $P(s \neq 0 \,|\, \theta_{s\neq 0})$ is the weight component. Note that all categories, except $\breve{s}_1 = 0$, have equal probability. We place a Beta prior, $p(\boldsymbol{\theta}_s \,|\, \frac{\alpha_{s=0}}{2}, \frac{\alpha_{s\neq 0}}{2})$, with hyperparameters $\alpha_{s=0}$ and $\alpha_{s\neq 0}$, over the mixture weights $\boldsymbol{\theta}_s = \{\theta_{s=0}, \theta_{s\neq 0}\}$ with $\theta_{s\neq 0} = 1 - \theta_{s=0}$. Marginalizing over $\theta_{s=0}$ yields,

$$
\begin{aligned}
P(\mathbf{S}) = \prod_{k=1}^{K} \int_0^1 \prod_{n=1}^{N_\mathrm{z}} \Big(& P(s_{n,k} = 0 \,|\, \theta_{s=0}) \delta(s_{n,k}) \\
& + P(s_{n,k} \neq 0 \,|\, \theta_{s=0})(1 - \delta(s_{n,k}) \,) \Big) \\
& \times p(\boldsymbol{\theta} \,|\, \alpha_{s=0}, \alpha_{s\neq 0}) \, \mathrm{d}\theta_{s=0} \\
\propto \prod_{k=1}^{K} & \, \mathrm{BetaBin}_{\mathbf{s}_k} \big(m_{s=0,k} + \alpha_{s=0}, m_{s\neq 0,k} + \alpha_{s\neq 0} \big) \;,
\end{aligned}
$$

where $\delta(s)$ returns one if $s = 0$ and zero otherwise, $m_{s=0,k}$ counts the zero elements in $\mathbf{s}_k$, and $m_{s\neq 0,k} = N_\mathrm{z} - m_{s=0,k}$.

3.5 Mixture of policies and action likelihood

Since we assume a finite set of actions, we consider a Categorical distribution for each mixture component,

$$
P(u \,|\, \boldsymbol{\phi}_k) = \mathrm{Cat}_u(\boldsymbol{\phi}_k) \;,
$$

such that the mixture model can be written as

$$
P(u \,|\, \boldsymbol{\Phi}, \mathbf{s}) = \frac{1}{Z_u} \sum_{k=1}^{K} s_k P(u \,|\, \boldsymbol{\phi}_k), \tag{6}
$$

with normalization constant $Z_u = \sum_{u\in\mathcal{U}} \sum_{k=1}^{K} s_k P(u \,|\, \boldsymbol{\phi}_k)$. The parameters of the policy of each mixture component follow Dirichlet distributions with identical hyperparameters, α_ϕ,

$$
p(\boldsymbol{\Phi}) = \prod_{k=1}^{K} \mathrm{Dir}_{\boldsymbol{\phi}_k}(\alpha_\phi, \ldots, \alpha_\phi) \;,
$$

where we assume independent policies. We consider the hyperparameter, α_ϕ, as a Gamma distributed variable with parameters $h^{(1)}_{\phi_A}$ and $h^{(2)}_{\phi_A}$.

As explained, we require a data set containing observed actions, u_n, $n = 1, \ldots, N_\mathrm{z}$. Assuming

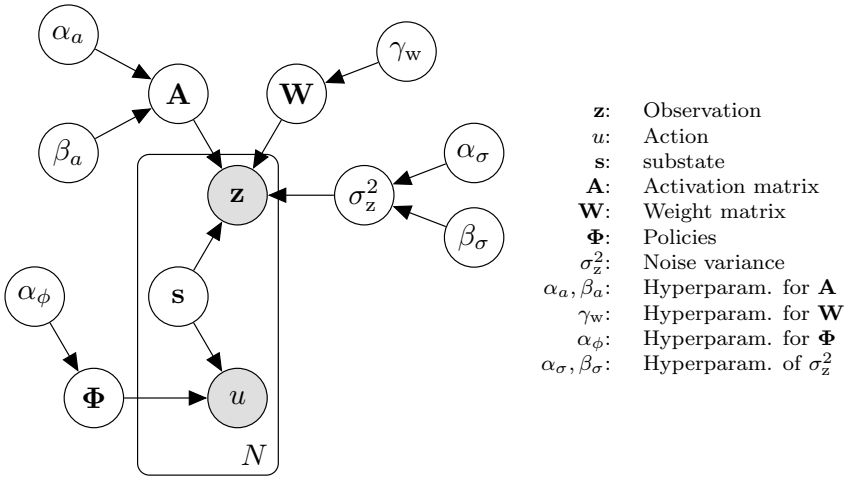

Figure 1: Graphical model of the feature-based decision-making model. Only the states, $\mathbf{z}_n$, and the actions, u_n, with $n = 1, \ldots, N_\mathrm{z}$, are observed. The other variables are latent and need to be inferred.

that the observed actions are independent, the likelihood is given as

$$P(\mathbf{u} \,|\, \mathbf{\Phi}, \mathbf{S}) = \prod_{n=1}^{N_\mathrm{z}} \frac{1}{Z_{u_n}} \sum_{k=1}^{K} s_k P(u_n \,|\, \boldsymbol{\phi}_k).$$

3.6 Joint posterior distribution

The full joint posterior distribution can be factorized as

$$\begin{aligned} p(\mathbf{W}, \mathbf{A}, \mathbf{S}, \mathbf{\Phi}, \sigma_\mathrm{z}^2, \gamma_\mathrm{w}, \alpha_a, \beta_a, \alpha_\sigma, \beta_\sigma, \alpha_\phi \,|\, \mathbf{Z}, \mathbf{u}) \propto \\ p(\mathbf{Z} \,|\, \mathbf{W}, \mathbf{A}, \mathbf{S}, \sigma_\mathrm{z}^2) P(\mathbf{u} \,|\, \mathbf{S}, \mathbf{\Phi}) p(\sigma_\mathrm{z}^2 \,|\, \alpha_\sigma, \beta_\sigma) \\ \times p(\mathbf{S}) p(\mathbf{W} \,|\, \gamma_\mathrm{w}) P(\mathbf{A} \,|\, \alpha_a, \beta_a) p(\mathbf{\Phi} \,|\, \alpha_\phi) \\ \times p(\gamma_\mathrm{w} | \alpha_\gamma, \beta_\gamma) p(\alpha_\sigma) p(\beta_\sigma) p(\alpha_a) p(\beta_a) p(\alpha_\phi). \end{aligned} \tag{7}$$

The conditional independences in this model are exploited in Section 4, where inference based on this model is explained. The structure of the posterior is illustrated as a graphical model in Fig. 1.

4 Inference

We consider the problem of learning the latent variables and the prediction of optimal actions for new observations as a Bayesian inference problem. Thus, we are interested in the joint posterior distribution in Eq. (7). Since the joint posterior is not directly tractable, we represent it by samples generated by means of Gibbs sampling. Therefore, in the first part of this section, we

derive the conditional distributions of the variables. After explaining the sampling scheme in the second part, in the third part, we detail how to predict actions for new observations using this model.

For convenience, we use the bar symbol $(-)$ in what follows to denote the set of conditional variables, i.e., all variables except the one that shall be sampled. The set of all latent variables is denoted by

$$\Omega = \{\mathbf{W}, \mathbf{A}, \mathbf{S}, \mathbf{\Phi}, \sigma_{\mathrm{z}}^2, \gamma_{\mathrm{w}}, \alpha_a, \beta_a, \alpha_\sigma, \beta_\sigma, \alpha_\phi\} \in \mathbf{\Omega},$$

where $\mathbf{\Omega}$ denotes the joint space of the latent variables.

4.1 Conditional distributions

4.1.1 Sampling the noise variance

The hyperpriors of α_σ and β_σ are conjugate to the prior of σ_{z}^2. Hence, the conditional $p(\sigma_{\mathrm{z}}^2 \,|\, -)$ is also Inverse-Gamma distributed,

$$p(\sigma_{\mathrm{z}}^2 \,|\, -) \propto \mathrm{IGa}_{\sigma_{\mathrm{z}}^2}\Bigg(\alpha_\sigma + \frac{N_{\mathrm{z}} D}{2} +, \beta_\sigma + \frac{1}{2}\sum_{n=1}^{N_{\mathrm{z}}}\sum_{d=1}^{D}\left(z_{n,d} - \sum_{k=1}^{K} s_{n,k} a_{k,d} w_{k,d}\right)^2\Bigg) .$$

We use a Metropolis-Hastings algorithm with a Gamma proposal distribution to generate samples of the hyperparameters, α_σ and β_σ.

4.1.2 Sampling the substates

Sampling the substates, $\mathbf{S}$, is simple thanks to the assumption of a finite space of the elements. The conditional consists of the state and action likelihoods as wells as the (conditional) prior,

$$P(s_{n,k} \,|\, -) \propto p(\mathbf{z}_n \,|\, \mathbf{W}, \mathbf{A}, \mathbf{s}_n, \sigma_{\mathrm{z}}^2) P(u_n \,|\, \mathbf{s}_n, \mathbf{\Phi}) P(s_{n,k} \,|\, \mathbf{S} \backslash s_{n,k}), \tag{8}$$

for all $s_{n,k} \in \mathcal{S}$, where $\mathbf{S}\backslash s_{n,k}$ denotes the elements of $\mathbf{S}$ without $s_{n,k}$. As the prior for the substates follows a Beta-Binomial distribution, the conditional is simply given as

$$P(s_{n,k} \,|\, \mathbf{S}\backslash s_{n,k}) \propto \begin{cases} m_{s=0} + \alpha_{s=0} & \text{for } s_{n,k} = 0 \\ m_{s\neq 0} + \alpha_{s\neq 0} & \text{for } s_{n,k} \neq 0 \end{cases} ,$$

where $m_{s=0}$ and $m_{s\neq 0}$ are defined in Section 3.4.

4.1.3 Sampling the feature weights

Since the likelihood is Gaussian and the weights are i.i.d. following an Exponential distribution, the conditional of the kth row of the weight matrix, $\mathbf{w}_k$, is a truncated Gaussian distribution,

$$\begin{aligned} p(\mathbf{w}_k \,|\, -) &\propto p(\mathbf{Z} \,|\, \mathbf{W}, \mathbf{A}, \mathbf{S}, \sigma_{\mathrm{z}}^2) \prod_{d=1}^{D} p(w_{k,d} \,|\, \gamma_{\mathrm{w}}) \\ &\propto \mathcal{TN}_{\mathbf{w}_k}\left(\boldsymbol{\mu}_{\mathbf{w}_k}, \boldsymbol{\Sigma}_{\mathbf{w}_k}\right) , \end{aligned}$$

with $\mathcal{TN}_{\mathbf{w}_k}\left(\boldsymbol{\mu}_{\mathbf{w}_k}, \boldsymbol{\Sigma}_{\mathbf{w}_k}\right)$ denoting a truncated Gaussian, where the elements of $\mathbf{w}_k$ are constraint to be positive-valued. The mean, $\boldsymbol{\mu}_{\mathbf{w}_k}$, and the covariance, $\boldsymbol{\Sigma}_{\mathbf{w}_k}$, are given as

$$\boldsymbol{\Sigma}_{\mathbf{w}_k}^{-1} = \frac{1}{\sigma_{\mathrm{z}}^2} \sum_{n=1}^{N_{\mathrm{z}}} s_{n,k}^2 \mathrm{diag}(\mathbf{a}_k) \tag{9}$$

$$\boldsymbol{\mu}_{\mathbf{w}_k} = \boldsymbol{\Sigma}_{\mathbf{w}_k}^{-1} \Big(\frac{1}{\sigma_{\mathrm{z}}^2} \sum_{n=1}^{N_{\mathrm{z}}} s_{n,k} \Big(\mathbf{z}_n - \sum_{\substack{k'=1 \\ k' \neq k}}^{K} s_{n,k'} \left(\mathbf{a}_{k'} \odot \mathbf{w}_{k'} \right) \Big) \odot \mathbf{a}_k - \mathbf{1}_D \frac{1}{\gamma_{\mathrm{w}}} \Big). \tag{10}$$

We use the algorithm presented in [7] to sample from a truncated Gaussian distribution. Note that we neglect the dependency on the activations in Eq. (9) since the covariance would be infinite for $a_{k,d} = 0$. This does not affect the sampling scheme, as the corresponding weight will be ignored due to $a_{k,d} = 0$ anyway. Sampling the hyperparameter, γ_{w}, is fairly easy since the conditional of γ_{w} is an Inverse-Gamma distribution,

$$p(\gamma_{\mathrm{w}} \,|\, -) \propto \mathrm{IGa}_{\gamma_{\mathrm{w}}} \left(\alpha_\gamma + \frac{KD}{2}, \beta_\gamma + \frac{1}{2} \sum_{k=1}^{K} \sum_{d=1}^{D} w_{k,d} \right) .$$

4.1.4 Sampling the feature activations

Sampling from the IBP consists of two steps: For each row, (i) the active columns are updated and then (ii) new features are proposed. In the first step, an element of the activation matrix, $a_{k,d}$, is set active with probability

$$P(a_{k,d} \,|\, -) \propto p(\mathbf{z}_d \,|\, \mathbf{S}\mathbf{f}_d, \sigma_{\mathrm{z}}^2) P(a_{k,d} \,|\, \mathbf{a}_{k \backslash d}), \tag{11}$$

with $\mathbf{a}_{k \backslash d}$ denoting the kth row of $\mathbf{A}$ without the dth element and $\mathbf{f}_d$ the dth column of $\mathbf{F}$. As $\mathbf{A}$ is assumed to follow an IBP, the conditional in Eq. (11) is [13, 14]

$$P(a_{k,d} = 1 \,|\, \mathbf{a}_{k \backslash d}) = \frac{m_{k \backslash d}}{D + \beta_a - 1},$$

where $m_{k \backslash d}$ is the sum over $\mathbf{a}_{k \backslash d}$.

In the second step, K^+ new features are proposed in a Metropolis step [10, 22]. The proposal

distribution, $q(\theta^+ \,|\, \theta)$, is independent of the previous sample, θ, as it consists of the priors for the substates, the feature weights, and the policies of the features,

$$q(\theta^+ \,|\, \theta) = q(\theta^+) = P(K^+ \,|\, -)p(\mathbf{W} \,|\, \gamma_\mathrm{w})P(\mathbf{S})p(\mathbf{\Phi} \,|\, \alpha_\phi) \tag{12}$$

with $\theta = \{\mathbf{W}, \mathbf{S}, \mathbf{A}, \mathbf{\Phi}\}$ and $\theta^+ = \{\mathbf{W}^+, \mathbf{S}^+, \mathbf{A}^+, \mathbf{\Phi}^+\}$, where $\mathbf{W}^+, \mathbf{S}^+, \mathbf{A}^+$, and $\mathbf{\Phi}^+$ denote the proposed feature weights, activations, coefficients, and policies. The probability of adding K^+ features is given as [13, 14]

$$P(K^+ \,|\, -) \sim \mathrm{Poisson}_{K^+}\left(\frac{\alpha_a \beta_a}{\beta_a + D - 1}\right).$$

As the proposal distribution is independent of the previous sample, the acceptance ratio, r, is equal to the likelihood ratio between the new and existing features [26],

$$r = \frac{P(\mathbf{Z} \,|\, \theta^+)}{P(\mathbf{Z} \,|\, \theta)}.$$

Since the IBP tends to mix slowly, we augment this ratio with probability P^+ of accepting a single new feature, resulting in a modified acceptance ratio which is derived in [22]. This increases the probability of proposing new features, leading to a faster convergence to the stationary distribution. The hyperparameters α_a and β_a are sampled as described in [22].

4.1.5 Sampling the policies

Sampling from the conditionals for the policies directly is difficult and would be computationally expensive due to the mixture model (Eq. (6)). An efficient approach is to introduce auxiliary variables, t_n, $n = 1, \ldots, N_\mathrm{z}$, for each observation, indicating from which policy the observed action, u_n, has been generated [25]. Given the indicators, the mixture components $\boldsymbol{\phi}_k$, $k = 1 \ldots, K$, in Eq. (6) become conditionally independent of the mixture weights, $\mathbf{s}_n$, $n = 1, \ldots, N_\mathrm{z}$, which makes sampling the components straightforward. The sampling algorithm thus consists of two steps.

First, the indicators are sampled according to

$$P(t_n = k \,|\, u_n, \mathbf{S}, \boldsymbol{\phi}) \propto s_k P(u_n \,|\, \boldsymbol{\phi}_k). \tag{13}$$

In order to approximate the conditional for the policies, we draw N_t indicator samples from Eq. (13). Drawing the samples is easy, since the indicators follow Categorical distributions with $t_n \in \{1, \ldots, K\}$.

Second, given the indicators, the parameters of the kth feature policy, $\boldsymbol{\phi}_k$, is sampled from a Dirichlet-Multinomial distribution,

$$p(\boldsymbol{\phi}_k \,|\, \mathbf{u}, \mathbf{t}) = \mathrm{DirMult}_{\boldsymbol{\phi}_k}(m_{\phi_k,1} + \alpha_\phi, \ldots, m_{\phi_k,N_\mathrm{u}} + \alpha_\phi)$$

where $m_{\phi_k,i}$, $i = 1, \ldots, N_\mathrm{u}$, counts the co-occurrences between the policy indicators, $\mathbf{t}$, and the actions, $u_n \in \mathcal{U}$.

The hyperparameter, α_ϕ, can be sampled in a Metropolis step, using a Gamma proposal distribution.

4.2 Sampling algorithm

The Gibbs sampler is initialized with only one feature and the variables are sampled from their prior distributions. After several iterations of the Gibbs sampler, samples from the target distribution are generated.

Note that there is the chance to generate new features in each iteration of the Gibbs sampler. Especially in scenarios with strong noise, different rows of the feature matrix may converge to similar realizations, increasing the number of features unnecessarily. We propose to merge features reducing the number of features, if they show a similarity larger than a prefixed threshold, T_corr, where we keep the activations of both features and average the policies. The similarity is measured by means of the estimated correlation between the feature samples.

A *maximum-a-posteriori* (MAP) estimator can be utilized, if we are interested in an estimate of the latent variables, Ω_MAP, containing, i.a., estimates of the features, $\hat{\mathbf{F}}_\mathrm{MAP}$ [35], and the policies, $\hat{\mathbf{\Phi}}_\mathrm{MAP}$. The MAP estimator can be approximated by choosing the sample with the highest posterior probability. The posterior probability is calculated according to Eq. (7). For the prediction of actions for new observations, we can also use a Minimum Mean Squared Error (MMSE) estimator which provides better generalization capabilities. This is detailed in the next section.

4.3 Prediction of actions

The proposed model can be used to learn the structure of the observed states as well as for the prediction of actions, given new observed states. For the prediction of an optimal action, $u^\star$, given a new observation, $\mathbf{z}^\star$, we can evaluate the posterior predictive distribution, giving rise to a MMSE estimator,

$$P(u^\star \,|\, \mathbf{z}^\star, \mathcal{D}) = \int_\mathcal{S} \int_\mathbf{\Omega} P(u^\star, \mathbf{s}^\star \,|\, \mathbf{z}^\star, \Omega) p(\Omega \,|\, \mathbf{z}^\star, \mathcal{D}) \,\mathrm{d}\Omega \,\mathrm{d}\mathbf{s}^\star, \tag{14}$$

where we exploit that $u^\star$ is independent of the data set containing the observations, $\mathcal{D}$, given Ω. Eq. (14) shows that Ω depends on the new observations, $\mathbf{z}^\star$. Thus, all variables would need to be inferred for each prediction, which would be computationally expensive. To remedy this issue, we assume that the observed data in $\mathcal{D}$ sufficiently represents the conditional distribution for Ω, such that we can ignore the dependency and simply infer Ω based on $\mathcal{D}$, leaving only $u^\star$ and $\mathbf{s}^\star$ to be inferred during prediction. The marginalization over Ω is approximated by Monte Carlo integration. For this, we need to draw samples of $\mathbf{s}^\star$ given the samples of Ω. The samples can be

Table 1: Parameter settings for the algorithm used in the simulation experiments

Parameter	Value	Meaning
$h^{(1)}_{\alpha_\sigma}$	1000	hyperparameters for σ_z
$h^{(2)}_{\alpha_\sigma}$	1	$h^{(1)}_{\alpha_\sigma}$ is chosen assuming a
$h^{(1)}_{\beta_\sigma}$	1	high Signal-To-Noise Ratio (SNR) scenario
$h^{(2)}_{\beta_\sigma}$	1	
$h^{(1)}_{\alpha_A}, h^{(2)}_{\alpha_A}$	1	hyperparameters for α_a
$h^{(1)}_{\beta_A}$	1	hyperparameter for β_a
$h^{(2)}_{\beta_A}$	10	hyperparameter for β_a
$\alpha_\gamma, \beta_\gamma$	1	hyperparameters for $\mathbf{W}$
$h^{(1)}_{\phi_A}, h^{(2)}_{\phi_A}$	1	hyperparameters for $\mathbf{\Phi}$
α_s, β_s	1	hyperparameters for $\mathbf{S}$
P^+	0.01	probability of accepting $K^+ = 1$ features
T_{corr}	0.9	threshold for merging similar features
N_{iter}	10000	number of iterations of the Gibbs Sampler
L	100	size of $\mathcal{S}$
N_t	1000	number of samples of the policy indicators

generated by drawing from the conditional in Eq. (8). Alternatively, we obtain a MAP estimate of $u^\star$ by maximizing the posterior predictive distribution with respect to $\mathbf{s}^\star$ given Ω_{MAP}.

Since, as justified in Section 2.3, we assume a deterministic policy, the optimal action, $u^\star_{\text{opt}}$, is the action that maximizes $P(u^\star \,|\, \mathbf{z}^\star, \mathcal{D})$. Thus, the posterior predictive distribution expresses our confidence about the actions and can be considered as the policy of the agent.

Note that, especially when we are interested in the prediction of actions, a modification on the model can help to increase the prediction accuracy dramatically. Especially with high-dimensional observation, the observation likelihood determines the posterior, where the effect of the action likelihood nearly vanishes, leading in the worst case to a neglect of the actions. A remedy consists in considering an action variable for each entry of the observation. However, we assume that these actions variables are identical. This increases the influence of the action log-likelihood by factor D, increasing the weight on the posterior significantly. The necessary modification in the inference algorithm are simple: First, for the conditional of Eq. (8), the action log-likelihood is multiplied by D. Second, the same modification is applied to the joint posterior distribution in Eq. (7). Sampling the policies is not affected as the actions for each substate are identical. We apply this modification on the model for the real data experiments in Section 6, as the observations in the experiments are high-dimensional.

5 Experimental Results

In order to demonstrate the performance of the proposed method, we consider simulations as well as real data experiments. In this section, we evaluate the performance of the inference

algorithm by simulating observations with different SNRs and different latent numbers of features. For this, we simulate the ground truth values by drawing samples from distributions which are similar to the prior distributions of the variables. This is detailed in the following. The true hyperparameters of the features weights are set to $\check{h}_{\alpha_\gamma} = \check{h}_{\beta_\gamma} = 100$. Thus, the true hyperparameter for the weights, $\check{\gamma}_\mathrm{w}$, as well as the true weights, $\check{\mathbf{W}}$, are sampled from their prior distributions. An element of the true feature activation matrix, $\check{\mathbf{A}}$, is activated with probability $P(\check{a}_{k,d} = 1) = 0.5$. The true substates, $\check{\mathbf{s}}_n$, are simulated by drawing samples from a Dirichlet distribution with parameters $\frac{1}{K}\mathbf{1}_K$, resulting in a peaked distribution. The true noise variance, $\check{\sigma}_\mathrm{z}^2$, is determined by the chosen SNR, where the signal power is estimated from $\check{\mathbf{X}} = \check{\mathbf{S}}\check{\mathbf{F}}$. The parameters of the ground truth policies, $\check{\mathbf{\Phi}}$, are chosen such that one action has high probability mass, reflecting deterministic policies, as justified in Section 2.3. In all simulations, we consider $N_\mathrm{z} = 100$ observations of dimensionality $D = 30$ with $N_\mathrm{u} = 4$ different actions. In order to evaluate the prediction performance of the model, we split the data set into a training and a test data set, leaving 80 observations for training and 20 for testing. We sweep the SNR from 10 dB to 30 dB with a step size of 5 dB and vary the number of features K within the set of $\{5, 7, 9, 12, 15, 18\}$. We set the hyperparameters of the algorithm according to Tab. 1.

We organize the evaluation in three parts. First, we investigate the performance of the estimation of the features, the feature coefficients, and the reconstruction of the states. Second, we compare the estimated policies to the true policies and evaluate the prediction in terms of the accuracy, A_u, which is the average rate of correctly predicted actions of the test data set. These estimates are obtained utilizing a MAP estimator. As for the prediction of actions we can also use the MMSE estimator, as detailed in Section 4.3, we also discuss the results obtained by this estimator. Third, we provide results for the inferred number of features and compare it with the true value.

5.1 Estimation of the features

The quality of the MAP estimates is evaluated in terms of the Mean Squared Error (MSE) for the elements of the features, $\mathbf{F}_\mathrm{MAP}$, the substates, $\mathbf{S}_\mathrm{MAP}$, and the reconstructions, $\mathbf{X}_\mathrm{MAP} = \mathbf{S}_\mathrm{MAP}\mathbf{F}_\mathrm{MAP}$. Further, we compute the RMSEs over 20 Monte Carlo runs, yielding RMSE measures for each estimated variable, $\epsilon_\mathbf{F}$, $\epsilon_\mathbf{S}$, and $\epsilon_\mathbf{X}$. The results are shown in Fig. 2(a)–(c).

We obtain good results with low errors especially for a small latent number of features, almost independent of the SNR. With increasing K, we observe a strong growth of the error. The error becomes even more significant in case of strong noise. It is remarkable that the error of the estimated feature values as well as the feature coefficients behave similarly. Due to the linear relation between the features and the substates, the errors of the reconstructions also grows with the number of features.

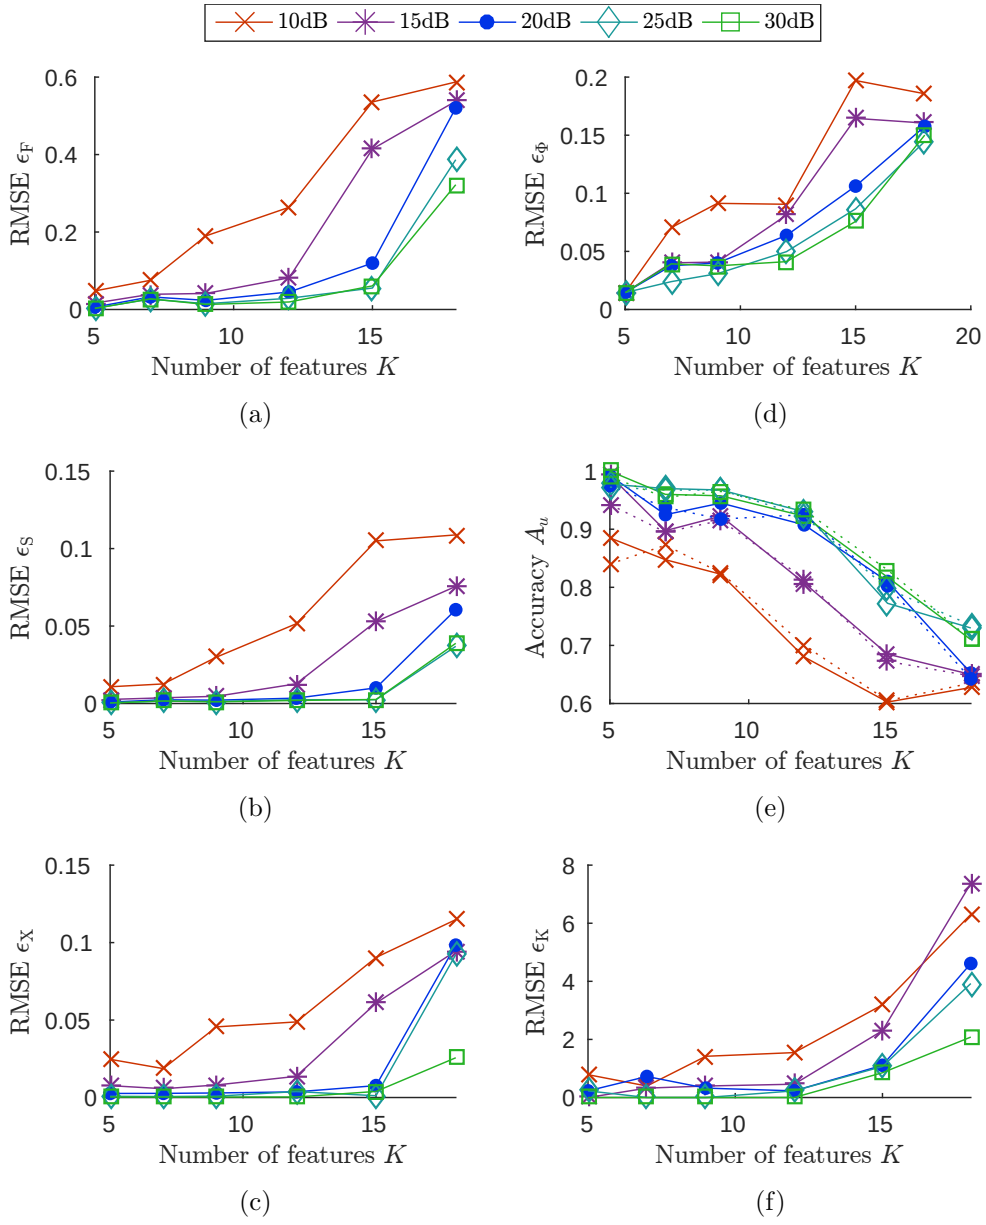

Figure 2: Results of the feature reconstruction. Shown are the RMSEs for (a) the features, (b) the substates, (c) the reconstructed observations, (d) the policies, (e) the accuracies of the predicted actions, and the (f) number of features. All variables are reliably inferred in case of moderate to high SNRs (SNR $>$ 15 dB). If the SNR drops below 15 dB, a reliable reconstruction of more than seven features cannot be guaranteed using the simulation settings.

5.2 Action prediction

We evaluate the correctness of the estimated policies in terms of the RMSE of the Mean Absolute Deviation (MAD) over 20 Monte Carlo runs. As shown in Fig. 2(d), the error slightly grows with an increasing number of features. Again, the SNR has, compared to the number of features, only little effect on the accuracy of the policies.

The accuracy of the action prediction using the MAP and MMSE estimators are depicted in Fig. 2(e). In case of few features ($K < 15$) and low noise (SNR > 20 dB), we obtain highly accurate predictions with over 90 % accuracy. Only for strong noise and many latent features the accuracy drops slightly below 70 %. This observation can be explained by the relation between the action and substates, which are challenging to infer in case of many features.

Since we assume in the simulation experiments that there is a fixed set of parameters that explains the observations, the MMSE estimator yields only minor improvements over the MAP estimator concerning the accuracy of predicted actions, as indicated by the dotted lines in Fig. 2(e).

5.3 Estimation of the number of features

Assuming that the observations have been generated by a fixed number of features, we evaluate how accurately the algorithm is able to infer this number. The results for the simulations are depicted in Fig. 2(f), showing the RMSE of the MAD over the 20 Monte Carlo runs. As can be observed, the MAP estimator is able to infer the correct number of features reliably, especially in case of few features. On the one hand, if the noise in the observations increases and the observations are based on many latent features, the error grows significantly. On the other hand, if the SNR is reasonably high, i.e., SNR > 15 dB, the results deviate on average by only four from the true number in case of 18 latent features. The error in estimating the number of features is probably due to the fact that some simulation examples can be explained by less features than used for their generation.

6 Real Data Experiments

We consider the problem of analyzing a driver's behavior, which is an important task for user-adaptive driver assistance systems [24, 47], in order to demonstrate the performance of the proposed model in a real-world scenario. For this, we observe the surrounding of the vehicle and the actions taken by the driver, aiming at learning what caused the driver to make the observed decisions. Using the proposed model, we can also predict which action the driver is likely to take given a certain situation. Thus, we also investigate the predictive performance of our approach by randomly creating training and test data sets. For this, we consider real data provided by the KITTI Vision Benchmark Suite [11] containing several challenges in urban driving. We use the data for the tracking challenge as it contains time-sequential LIDAR measurements of different situations in public road traffic. A schematic plot of the setup is illustrated in Fig. 3.

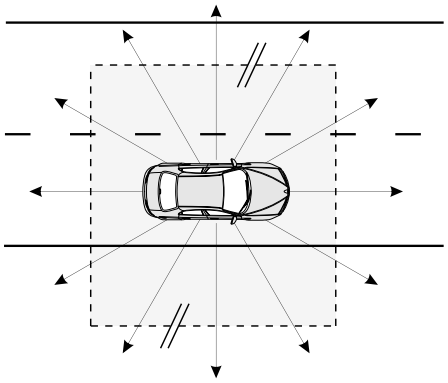

Figure 3: Illustration of the setup for the real data experiments. During pre-processing, we map the LIDAR measurement (indicated by the arrows) to an occupancy grid which is illustrated by the gray area surrounding the vehicle.

We consider Scene 11 and 20 of the benchmark suite, which are detailed below in Section 6.1 and Section 6.2. Before running the proposed inference algorithm, we pre-process the data as follows. First, we apply a thresholding on the height values of the measurements such that we keep only samples above ground level, which is roughly 1.5 m below the LIDAR. Afterwards, we discretize the measurements to obtain positive-valued occupancy grids, where the values of the grid elements refer to the maximum measured height.

As observation in the nth time frame, we consider the joint occupancy grid of the current and the previous frame. This is required to implicitly include velocity information, which enables to decide, e.g., if the host vehicle is faster or slower than the vehicle in front. Thus, the minimum speed, $v_{\min}$, between the host vehicle, H, and an obstacle, O, that can be resolved depends on the resolution, R, of the grid as follows,

$$R \leq |\, d_n - d_{n-1} \,| = (t_{\mathrm{F},n} - t_{\mathrm{F},n-1})\, |\, v_{\mathrm{H}} - v_{\mathrm{O}} \,| = \frac{1}{f_{\mathrm{s}}} v_{\min},$$

where d_n denotes the distance between the host vehicle and the obstacle in the nth time frame. The time stamp is denoted by $t_{\mathrm{F},n}$ and f_{s} is the frequency at which the measurements are sampled (in the KITTI data set, $f_{\mathrm{s}} = 10$ Hz).

We consider the environment of the vehicle in the range from -3 m to 3 m in the lateral direction. This covers the lane of the vehicle plus parts of (if existent) neighboring lanes. In Scene 11, the longitudinal range is limited between -10 m and 30 m and for Scene 20 between -40 m and 40 m. The range in longitudinal direction gives a time window of more than 2 s to react to the observed situation when driving at 50 km/h, which is the speed limit in German cities.

For both scenes, we choose a grid size of 21×65 pixels, resulting in a spatial resolution in lateral direction of $R_{\mathrm{lat}} \approx 0.3$ m and in longitudinal direction of $R_{\mathrm{long}} \approx 0.6$ m for Scene 11 (and $R_{\mathrm{long}} \approx 1.2$ m for Scene 20). Thus, the minimum velocities that can be detected are

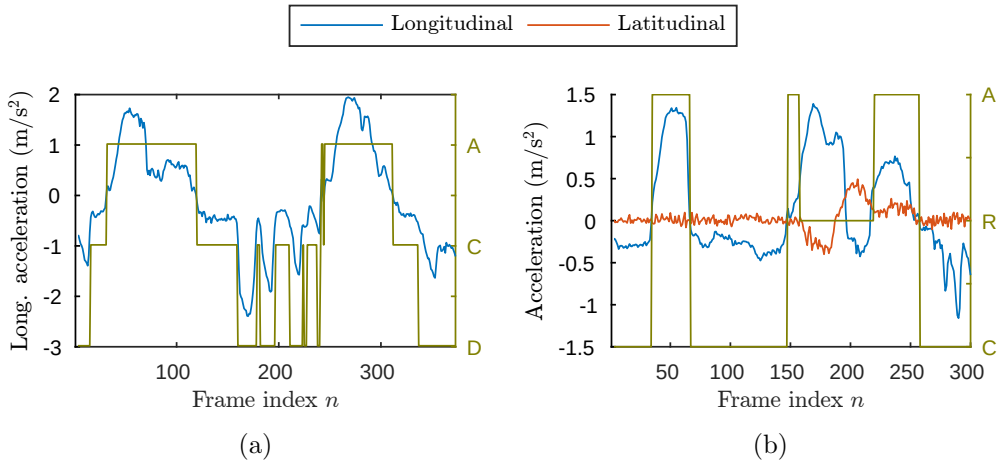

Figure 4: Ground truth labels of (a) Scene 11 and (b) Scene 20. The actions *acceleration* (A), *lane change (right)* (R), *deceleration* (D), and *moving at constant speed* (C) are chosen according to the acceleration of the vehicle.

$v_{\text{lat,min}} > 3$ m/s in the lateral direction and $v_{\text{long,min}} > 6$ m/s in the longitudinal direction (and $v_{\text{long,min}} > 12$ m/s for Scene 20).

We labeled the scenes to obtain the observed actions by means of the measured acceleration of the host vehicle, obtained from the onboard inertial measurement unit. The ground truth for Scene 11 is depicted in Fig. 4(a) and for Scene 20 in Fig. 4(b). The set of actions, $\mathcal{U}$, we consider in this experiment is comprised of *acceleration* (A), *deceleration* (D), *lane change (right)* (R), and *moving at constant speed* (C). However, due to the short duration of the sequences, usually fewer actions are observed in each scene.

We use the same settings for the hyperparameters as in Tab. 1. Only one of the two hyperparameter for the IBP, $h^{(2)}_{\alpha_A}$, is modified to reduce the number of expected features, $h^{(2)}_{\alpha_A} = 10$. Further, we perform 500 iterations of the Gibbs sampler. In both data sets, after a mere 100 iterations, the sampler converges to the stationary distribution and produces samples from the target distribution.

6.1 Scene 11 - Traffic jam

Scene 11 of the KITTI data set shows an urban scenario, in which the driver follows another vehicle. The vehicle in front cannot be overtaken due to a single-lane road. As shown in Fig. 4(a), the driver first accelerates. After a few seconds, the driver has to decelerate until the car stops due to a halt of the preceding car. When the vehicle in front starts moving again, the host vehicle accelerates.

Applying the proposed algorithm to the observations results in 31 features using the MAP estimator. For the analysis of the features, we only visualize three of the most relevant features, which are selected according to the confidence in the corresponding feature policy.

Considering the current and the previous frames as inputs, we obtain feature estimates of consecutive frames. Observing the differences between the frames can be difficult as the features are likely to be highly similar due to the low temporal difference. Therefore, we show the feature of the current time frame and the difference of the features. A difference plot can be interpreted as the temporal gradient of two consecutive frames. Thus, from the depicted patterns we can draw conclusions about the relative speed of the obstacles. For instance, if the difference plot shows negative values left of positive values (red-blue pattern), the obstacle is faster than the host vehicle. In contrast, a blue-red pattern indicates a slower vehicle. The width of the bars reflect the magnitude of the relative velocities.

The features are illustrated in Fig. 5. In the figures, the x-axis corresponds to the longitudinal and the y-axis to the lateral direction. Shown is the top-view on the host vehicle, which is indicated by the red cross. The intensity reflects the measured heights at each pixel. As explained, the difference features are obtained by subtracting the features of the previous from the current time frame. Features A1 and A2 clearly show that the driver accelerates as long as the preceding car is significantly faster. A3, however, indicates acceleration though the preceding car decelerates. Since the car seems to be sufficiently far away, the driver has not yet decided to reduce the speed. Features D1 and D2 explain the deceleration of the driver with a slower car in front. In contrast, D3 shows that the preceding car accelerates. However, due to the low distance between both cars, the driver decides to decelerate. Features C1 and C2 show a vehicle in front of the driver's car. The difference images reveal only little differences between the velocities of both vehicles, such that the driver maintains the current velocity. C3 represents an empty road, where, again, the driver does not need to adapt the speed of his car.

In order to evaluate the prediction performance, as in the simulations, 20% of the observations are used for testing while the rest is used for inferring the structure. Using the MAP estimator results in an accuracy of 74.32 %. The confusion matrix in Tab. 2 shows that, most notably, in some cases moving at constant speed and acceleration are confused. Further, deceleration is in few cases misclassified as acceleration.

Quantifying the results of the state reconstruction is difficult, as a ground truth is not available. To provide at least an intuition about the quality of the reconstructions, we compute the MSE between the reconstructed states, based on the estimated substates and features, and the observations. This results in a MSE of approx. 0.0329. As the values of the observations are in a similar range as in the simulations, this result indicates low errors, resembling the results of the simulations. Still, the number has to be taken with care, as we compare the denoised reconstruction with noisy observation.

6.2 Scene 20 - Lane change

Scene 20 shows a lane change maneuver on a two-lane road. As can be observed from the acceleration signals depicted in Fig. 4(b), the driver accelerates three times, depending on the current traffic situation. The lane change takes place from time frame 158 to 220.

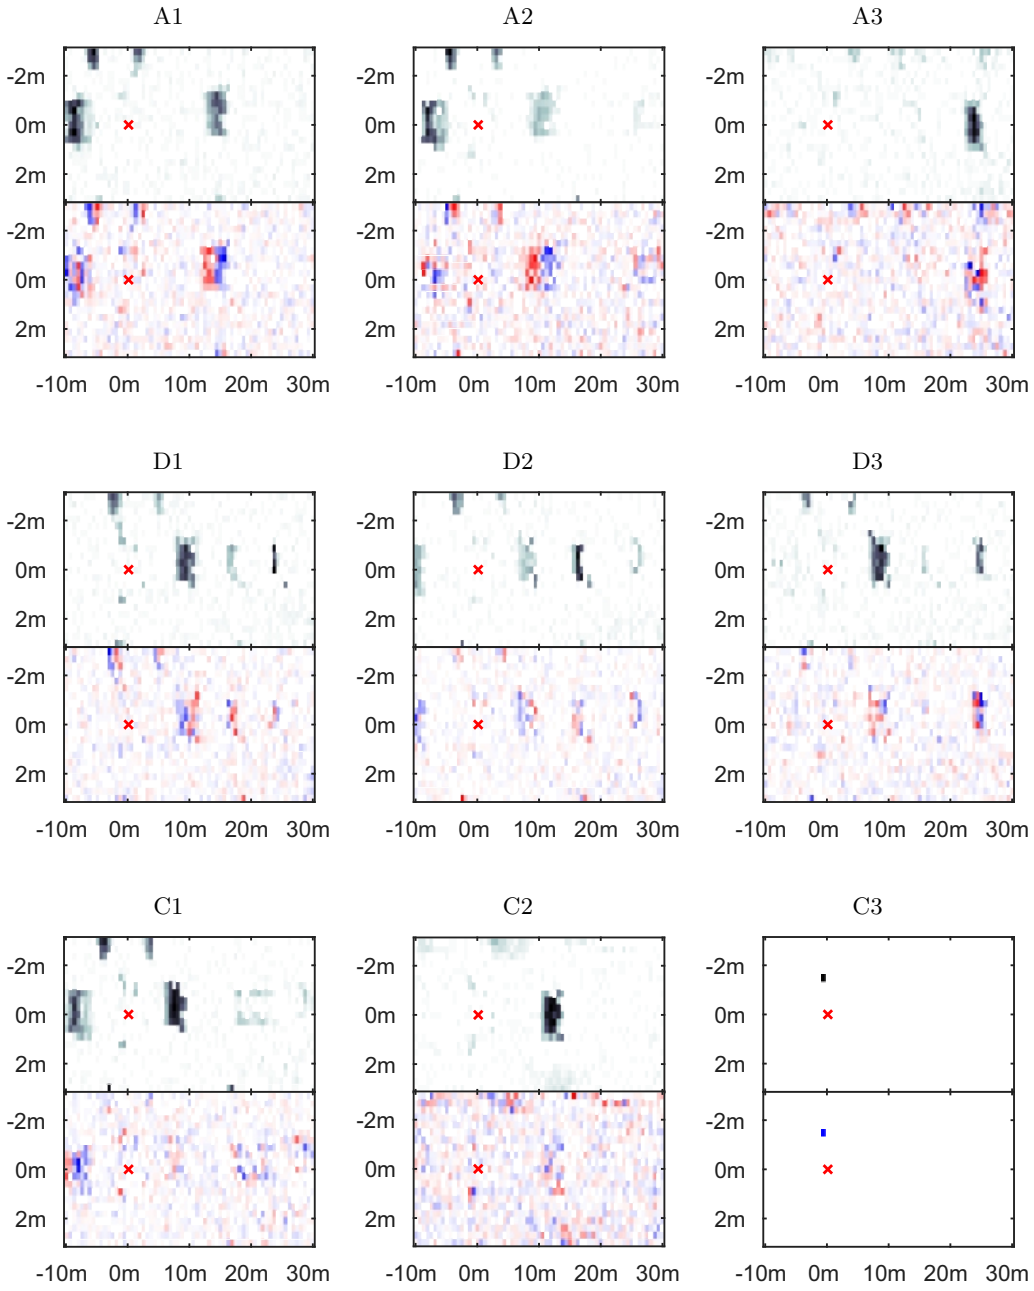

Figure 5: Features estimated for KITTI Scene 11. The host vehicle is indicated by the red cross. Depicted are the features (top) and the difference plots (bottom), where blue indicates positive and red negative values. First row: action A (acceleration). Features A1 and A2 show an obstacle in front of the host vehicle, which is moving significantly faster as indicated by the difference plots. A3 shows an obstacle that slows down. As it is sufficiently far away, the driver still accelerates. Second row: action D (deceleration). Features D1 and D2 show vehicles in front of the host car, which are moving at a slower speed. D3 shows that the preceding car accelerates. As the host vehicle is close to the obstacle, the driver decelerates. Third row: action C (moving at constant speed). Features C1 and C2 show preceding vehicles. As the difference plots do not indicate significant speed differences, the preceding and the host car travel at a similar speed, such that the driver does not need to adapt the speed. C3 shows an empty road without any obstacles.

Table 2: Confusion matrix for the hold-out data set of Scene 11. The overall accuracy is 74.32%, where 31 features have been inferred.

	Prediction		
Ground truth	Const. speed	Deceleration	Acceleration
Const. speed	**12**	2	4
Deceleration	1	**15**	6
Acceleration	5	1	**28**

Table 3: Confusion matrix for the hold-out data set of Scene 20. The overall accuracy is 73.33%, where 17 features have been inferred.

	Prediction		
Ground truth	Acceleration	Lane change (right)	Const. speed
Acceleration	**3**	3	7
Lane change (right)	0	**29**	0
Const. speed	0	6	**12**

Applying the proposed algorithm reveals 17 features. Using these features to predict the actions of the test data set yields an accuracy of 73.33%. The confusion matrix in Tab. 3 shows that lane changes are reliably predicted. Acceleration maneuvers are in some cases misinterpreted as moving at constant speed or as a lane change. A third of the moving at constant speed observations are misclassified as lane changes. Comparing the reconstructed states with the noisy observation yields a MSE of 0.0027.

For illustration of the inferred features, Fig. 6 shows three out of the 17 features, each indicating a different action. The first feature (a) shows a lane change (right). As the driver is already performing the maneuver, the scene is rotated. The dark areas in the lower left corner represent vehicles behind the host car. The second feature (b) contains a vehicle behind the driver's car and another vehicle on the left lane. Due to the high traffic density, the driver does not accelerate. In contrast, the third feature (c) shows a completely empty road, motivating the driver to accelerate the vehicle.

7 Discussion

As shown in the simulation experiments, the algorithm based on the proposed model is able to reliably infer the number of features, the features and the policies. In case of strong noise, the algorithm finds several, almost equally probable explanations for the observations, resulting in variations of the MAP estimate. Especially when the number of features is high compared to the dimensionality of the observations, inferring the correct number of features can be challenging.

The real data experiments show that the model is able to provide deeper insights into the observations which may yield new conclusions about the observed behavior. As explained, in high-dimensional observations, the observation likelihood is likely to dominate the posterior leading to only little influence of the action likelihood and, hence, poor prediction performance.

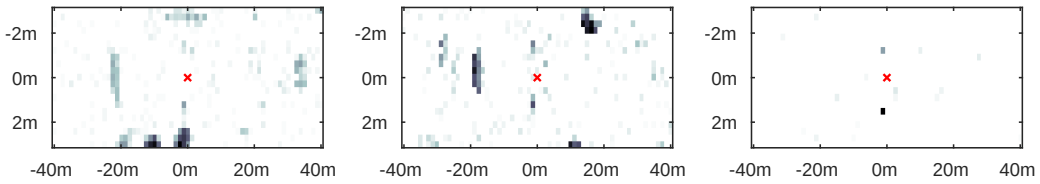

Figure 6: Features estimated from KITTI Scene 20. Shown are 3 out of 17 inferred features and the host vehicle (red cross). The first feature (a) indicates a lane change (right). The scene is rotated, showing two vehicles behind the host vehicle in the lower left corner. The second feature (b) represents moving at constant speed, most likely due to high traffic density (vehicle behind the host and on the left lane). The third feature (c) indicates acceleration and shows a free lane.

As proposed in Section 4.3, reweighting the action likelihood, assuming an action variable for each entry of the observation, yields a significant performance increase, while the observed states can still be reliably reconstructed.

An advantage of the proposed generative model is that it can be modified and extended easily. This is helpful especially for a different assumption on the feature weights. For example, if real-valued features are expected, the corresponding prior can be changed to a Gaussian distribution. Of course, inference has to be adapted accordingly. Further, one can easily extend the proposed model to a semi-supervised learning approach, in which we add variables for states where the actions are not observed. Thus, the state representations can be learned from even more data resulting in more accurate estimates.

8 Conclusion

We have presented a feature-based framework for learning from demonstrations that allows us to reason about the observed behavior and to predict actions for new states. A key assumption in the proposed model is that the observations are composed of latent features. Each feature imposes its own policy and contributes to the decision of the agent. To learn the structure of the behavior and to predict actions, we have considered a Bayesian nonparametric approach based on the Indian Buffet Process, which allows to infer the number of features and the features itself from the observed data. By means of this model, we are able to obtain a deeper understanding of the observed behavior as the features and their policies allow to reason about the observed decisions. The simulations show that the developed algorithm performs well. Only in scenarios with strong noise and many latent features, inference becomes challenging. Further, we have considered the task of learning a driver's behavior. For this, we have applied our algorithm to real data obtained from the KITTI benchmark suite. The results reveal under which conditions the driver takes the observed actions. Further, prediction on a hold-out data set demonstrates that actions can be predicted reliably.

References

[1] B. D. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.

[2] C. G. Atkeson and S. Schaal. Robot learning from demonstration. In *Proceedings of the 14th International Conference on Machine Learning*, Nashville, TN, USA, June 1997.

[3] R. Bellman. *Dynamic Programming*. Princeton University Press, Princeton, NJ, USA, 1st edition, 1957.

[4] C. Boutilier, R. Dearden, and M. Goldszmidt. Exploiting structure in policy construction. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 1104–1111, Montreal, Canada, Aug. 1995.

[5] S. J. Bradtke and A. G. Barto. Linear least-squares algorithms for temporal difference learning. *Machine Learning*, 22(1-3):33–57, 1996.

[6] J. Choi and K.-E. Kim. Nonparametric Bayesian inverse reinforcement learning for multiple reward functions. In *Advances in Neural Information Processing Systems 25*, pages 305–313, 2012.

[7] N. Chopin. Fast simulation of truncated Gaussian distributions. *Statistics and Computing*, 21(2):275–288, 2010.

[8] M. Daswani, P. Sunehag, and M. Hutter. Feature reinforcement learning: State of the art. In *AAAI-14 Workshop on Sequential Decision-making with Big Data*. Association for the Advancement of Artificial Intelligence, 2014.

[9] T. Degris, O. Sigaud, and P.-H. Wuillemin. Learning the structure of factored Markov decision processes in reinforcement learning problems. In *Proceedings of the 23rd International Conference on Machine learning*, pages 257–264, Pittsburgh, PA, USA, June 2006.

[10] F. Doshi-Velez and Z. Ghahramani. Accelerated sampling for the Indian buffet process. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 273–280, New York, NY, USA, June 2009.

[11] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *Proceedings of the 25th IEEE International Conference on Computer Vision and Pattern Recognition*, Tsukuba, Japan, Nov. 2012.

[12] S. J. Gershman, K. A. Norman, and Y. Niv. Discovering latent causes in reinforcement learning. *Current Opinion in Behavioral Sciences*, 5:43 – 50, 2015.

[13] Z. Ghahramani and T. L. Griffiths. Infinite latent feature models and the Indian buffet process. In *Advances in Neural Information Processing Systems 18*, pages 475–482, 2005.

[14] Z. Ghahramani, P. Sollich, and T. L. Griffiths. Bayesian nonparametric latent feature models. In *Bayesian Statistics 8*. Oxford University Press, Oxford, 2007.

[15] T. L. Griffiths and Z. Ghahramani. The Indian buffet process: An introduction and review. *Journal of Machine Learning Research*, 12:1185–1224, 2011.

[16] I. Guyon and A. Elisseeff. An introduction to variable and feature selection. *Journal of Machine Learning Research*, 3:1157–1182, 2003.

[17] J. Hahn and A. M. Zoubir. Inverse reinforcement learning using expectation maximization in mixture models. In *Proceedings of the 40th IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3721–3725, Brisbane, Australia, Apr. 2015.

[18] M. Hutter. Feature reinforcement learning: Part I. Unstructured MDPs. *Journal of Artificial General Intelligence*, 1(1):3–24, 2009.

[19] H. Ishwaran and J. S. Rao. Spike and slab variable selection: Frequentist and Bayesian strategies. *Annals of Statistics*, pages 730–773, 2005.

[20] I. T. Jolliffe. *Principal component analysis*. Springer, New York, 2nd edition, 2002.

[21] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1):99–134, 1998.

[22] D. Knowles and Z. Ghahramani. Nonparametric Bayesian sparse factor models with application to gene expression modeling. *Annals of Applied Statistics*, 5(2B):1534–1552, June 2011.

[23] D. D. Lee and H. S. Seung. Algorithms for non-negative matrix factorization. In *Advances in Neural Information Processing Systems 14*, pages 556–562, 2001.

[24] J. C. McCall and M. M. Trivedi. Driver behavior and situation aware brake assistance for intelligent vehicles. *Proceedings of the IEEE*, 95(2):374–387, Feb 2007.

[25] G. McLachlan and D. Peel. *Finite Mixture Models*. John Wiley & Sons, New York, 2000.

[26] E. Meeds, Z. Ghahramani, R. M. Neal, and S. T. Roweis. Modeling dyadic data with binary latent factors. In *Advances in Neural Information Processing Systems 19*, pages 977–984, 2006.

[27] B. Michini and J. P. How. Bayesian nonparametric inverse reinforcement learning. In *Proceedings of the European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 148–163, Bristol, UK, Sept. 2012.

[28] T. J. Mitchell and J. J. Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988.

[29] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[30] A. Y. Ng and S. J. Russell. Algorithms for inverse reinforcement learning. In *Proceedings of the 17th International Conference on Machine Learning*, pages 663–670, San Francisco, CA, USA, June 2000.

[31] T. Nguyen, Z. Li, T. Silander, and T. Y. Leong. Online feature selection for model-based reinforcement learning. In *Proceedings of the 30th International Conference on Machine Learning*, pages 498–506, Atlanta, GA, USA, June 2013.

[32] S. Parisi, M. Pirotta, N. Smacchia, L. Bascetta, and M. Restelli. Policy gradient approaches for multi-objective sequential decision making. In *Proceedings of the International Joint Conference on Neural Networks*, pages 2323–2330, Beijing, China, July 2014.

[33] D. A. Pomerleau. Efficient training of artificial neural networks for autonomous navigation. *Neural Computation*, 3(1):88–97, 1991.

[34] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., New York, NY, USA, 1st edition, 1994.

[35] P. Rai and H. Daume. Beam search based map estimates for the Indian buffet process. In *Proceedings of the 28th International Conference on Machine Learning*, pages 705–712, Bellevue, WA, USA, June 2011.

[36] N. Ratliff, D. Bradley, J. A. Bagnell, and J. Chestnutt. Boosting structured prediction for imitation learning. In *Advances in Neural Information Processing Systems 20*, 2007.

[37] N. D. Ratliff, J. A. Bagnell, and M. A. Zinkevich. Maximum margin planning. In *Proceedings of the 23rd International Conference on Machine Learning*, pages 729–736, New York, NY, USA, June 2006.

[38] M. Riedmiller. Neural fitted Q iteration – first experiences with a data efficient neural reinforcement learning method. In *Proceedings of the 16th European Conference on Machine Learning*, pages 317–328, Porto, Portugal, Oct. 2005.

[39] S. Ross, G. J. Gordon, and D. Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics*, pages 627–635, Ft. Lauderdale, FL, USA, Apr. 2011.

[40] G. Shani, D. Heckerman, and R. I. Brafman. An MDP-based recommender system. *Journal of Machine Learning Research*, 6:1265–1295, 2005.

[41] C. R. Shelton. *Importance sampling for reinforcement learning with multiple objectives*. PhD thesis, Massachusetts Institute of Technology, 2001.

[42] R. D. Smallwood and E. J. Sondik. The optimal control of partially observable Markov processes over a finite horizon. *Operations Research*, 21(5):1071–1088, 1973.

[43] A. Šošić, A. M. Zoubir, and H. Koeppl. A Bayesian approach to policy recognition and state representation learning. *arXiv:1605.01278*, 2016.

[44] A. Sosic, A. M. Zoubir, and H. Koeppl. Policy recognition via expectation maximization. In *Proceedings of the 41st IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4801–4805, Mar. 2016.

[45] A. Surana and K. Srivastava. Bayesian nonparametric inverse reinforcement learning for switched Markov decision processes. In *Proceedings of the 13th International Conference on Machine Learning and Applications*, pages 47–54, Dec. 2014.

[46] P. Vamplew, R. Dazeley, E. Barker, and A. Kelarev. Constructing stochastic mixture policies for episodic multiobjective reinforcement learning tasks. In *Proceedings of the 22nd Australasian Joint Conference on Advances in Artificial Intelligence*, pages 340–349, Melbourne, Australia, Dec. 2009.

[47] J. Wang, C. Yu, S. E. Li, and L. Wang. A forward collision warning algorithm with adaptation to driver behaviors. *IEEE Transactions on Intelligent Transportation Systems*, 17(4):1157–1167, Apr. 2016.

[48] C. J. C. H. Watkins. *Learning from Delayed Rewards*. PhD thesis, Cambridge University, Cambridge, England., 1989.

[49] M. Wulfmeier, P. Ondruska, and I. Posner. Deep inverse reinforcement learning. In *NIPS Deep Reinforcement Learning Workshop*, Montréal, Canada, Dec. 2015.